

2

AEROSPACE REPORT NO.
TR-93(3411)-1

AD-A259 891



A Numerical Technique for the Hierarchical Evaluation of Large, Closed Fault-Tolerant Systems

Prepared by

D. A. LEE, JR. and G. C. GILLEY
Engineering and Technology Group
The Aerospace Corporation

J. ABRAHAM
Electrical and Computer Engineering Department
University of Texas, Austin

D. RENNELS
Computer Science Department
University of California, Los Angeles

1 October 1992

DTIC
SELECTED
FEB 04 1993
S B D

Prepared for

SPACE AND MISSILE SYSTEMS CENTER
AIR FORCE MATERIEL COMMAND
Los Angeles Air Force Base
P. O. Box 92960
Los Angeles, CA 90009-2960

Programs Group

THE AEROSPACE CORPORATION
El Segundo, California

93

2

3

004

93-01996



3901

APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION UNLIMITED

This final report was submitted by The Aerospace Corporation, El Segundo, CA 90245-4691, under Contract No. F04701-88-C-0089 with the Space Systems Division, P. O. Box 92960, Los Angeles, CA 90009-2960. It was reviewed and approved for The Aerospace Corporation by H. J. Wertz, Principal Director, Computer Engineering Subdivision, Computer Systems Division, Engineering and Technology Group. The program officer is Mario Miranda, SSD/MBT.

This report has been reviewed by the Public Affairs Office (PAS) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nations.

This technical report has been reviewed and is approved for publication. Publication of this report does not constitute Air Force approval of the report's findings or conclusions. It is published only for the exchange and stimulation of ideas.

FOR THE COMMANDER



Mario N. Miranda
Follow-on Early Warning System
Test and Operations Division

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE		Approved for public release; distribution unlimited	
4. PERFORMING ORGANIZATION REPORT NUMBER(S) TR-93(3411)-1		5. MONITORING ORGANIZATION REPORT NUMBER(S) SMC-TR-92-53	
6a. NAME OF PERFORMING ORGANIZATION The Aerospace Corporation	6b. OFFICE SYMBOL (if applicable)	7a. NAME OF MONITORING ORGANIZATION Space and Missile Systems Center Air Force Materiel Command	
6c. ADDRESS (City, State, and ZIP Code) 2350 E. El Segundo Blvd. El Segundo, CA 90245-4691		7b. ADDRESS (City, State, and ZIP Code) Los Angeles Air Force Base P. O. Box 92960 Los Angeles, CA 90009-2960	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Space and Missile Systems Center Air Force Materiel Command	8b. OFFICE SYMBOL (if applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER	
8c. ADDRESS (City, State, and ZIP Code)		10. SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO.	PROJECT NO.
		TASK NO.	WORK UNIT ACCESSION NO.
11. TITLE (Include Security Classification) A Numerical Technique for the Hierarchical Evaluation of Large, Closed Fault-Tolerant Systems			
12. PERSONAL AUTHOR(S) D. A. Lee, Jr; G. C. Gilley; J. Abraham; D. Rennels			
13a. TYPE OF REPORT TR	13b. TIME COVERED FROM Aug 1988 TO Aug 1990	14. DATE OF REPORT (Year, Month, Day) 1 October 1992	15. PAGE COUNT 38
16. SUPPLEMENTARY NOTATION			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	
19. ABSTRACT (Continue on reverse if necessary and identify by block number)			
<p>This report describes a new technique for predicting the reliability of large, closed, fault-tolerant systems that overcomes significant shortfalls inherent in the reliability prediction tools commonly used today by systems developers. This new technique deals effectively with the problems of large state space and the introduction of coverages at all levels in the design hierarchy. A preliminary version of a tool using this technique has been developed that runs efficiently on a personal computer, in terms of both time and required resources. This achievement is made possible through the use of semi-Markov models to model each level in the hierarchy, and an innovative numerical technique which combines the models from any given level for use at the next higher level. The accuracy of the tool has been validated through comparative analytical calculations and Monte Carlo simulations.</p>			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL Mario Miranda		22b. TELEPHONE (Include Area Code) (310) 363-0958	22c. OFFICE SYMBOL SMC/MBT

PREFACE

We wish to thank Venu Shamapant for performing the Monte Carlo simulations, Mark Joseph and Joe Bannister for their useful comments, Patrick Kennedy for his assistance in the analytical calculation of the example triad's reliability, and the developers of CRAFTS for providing the use of their tool for this work.

DTIC QUALITY INSPECTED 3

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution	
Availability Codes	
Dist	Avail and/or Special
A-1	

CONTENTS

I. INTRODUCTION	7
II. NUMERICAL TECHNIQUE FOR SOLVING A HIERARCHY OF SEMI-MARKOV MODELS	11
A. Description of the Technique	11
B. Current Limitations of the Technique	15
III. IMPORTANCE OF COVERAGE IN THE DESIGN HIERARCHY	17
IV. PREDICTING THE RELIABILITY OF A SYSTEM DESIGN	23
VI. CONCLUSIONS AND FUTURE WORK	31
BIBLIOGRAPHY	33
APPENDIX	35

TABLES

1. Tool Characteristics	8
2. Triad System Reliability, Coverage = 1	18
3. Triad System Reliability, Coverage = 0.99	18
4. Triad System Reliability, Coverage = 0.95	19
5. Triad System Reliability, Coverage = 0.90	19
6. Ratio of Model Unreliability for Triad	21
7. Multicomputer Reliability	26
8. Ratio of Model Unreliability for Multicomputer	27
9. Multicomputer Reliability, Coverage = 1	27
10. Multicomputer Reliability, Coverage = 0.99	28
11. Multicomputer Reliability, Coverage = 0.95	28
12. Multicomputer Reliability, Coverage = 0.90	29

FIGURES

1. Modeling Approach for Three Design Levels	13
2. Fault Tree for Three-Design-Level Approach	14
3. Organization of One Computer in SNARC Modeling	24
4. Fault Tree for Computer in SNARC Modeling	25

I. INTRODUCTION

This report describes a new technique for predicting the reliability of large, closed, fault-tolerant systems that overcomes significant shortfalls inherent in the reliability prediction tools commonly used today by system developers. This new technique deals effectively with the problems of large state spaces and the introduction of coverages at all levels in the design hierarchy. A preliminary version of a tool using this technique has been developed that runs efficiently on a personal computer (PC), in terms of both time and required resources. This achievement is made possible through the use of semi-Markov models to model each level in the hierarchy, and an innovative numerical technique which combines the models from any given level for use at the next higher level. The accuracy of the tool has been validated through comparative analytical calculations and Monte Carlo simulations. Once the development of the tool is completed, the tool could be inexpensively distributed by the government. Table I lists the characteristics of several representative tools.

Large state spaces pose a serious problem with reliability prediction tools that use Markov models. They occur because the number of states in the system model increases geometrically with the number of components modeled. As a result, the models can be so large that either the capability of the tool or the capacity of the computing resource hosting the tool is exceeded. CARE III [Bavu 84a, Bavu 84b] and HARP [Duga 86, Triv 85] suffer from this problem. Mathematical techniques have been introduced to ease the effects of this problem. However, models of current satellite systems typically can be so large that even these techniques do not help. The new tool introduced in this report avoids the large state space problem altogether by using the design's hierarchical organization to reduce the number of states in the model. Each design level is converted into separate semi-Markov models, and a numerical

Table 1. Tool Characteristics

Tools	Problems with Large State Spaces	Include Coverages Throughout the Design Hierarchy	Proprietary	Cost	PC Version that Solves Large Models	Widely Used by System Developers
CARE III	Yes	No	No	\$5,371 ¹	No	Yes
HARP ²	Yes	No	No	Free	No ³	Yes
CRAFTS ⁴	No	No	Yes	\$25,000	No	No
SHARPE ⁵	No	Yes	Yes	\$10,000	No	No
SNARC	No	Yes	No	Free	Yes	No

¹The cost includes \$1,278 for the CARE III User Friendly Interface Program and User Friendly Interface documentation.

²Version 1.1 was used in this effort.

³A PC version of Harp exists; however, the PC version is unable to solve large models.

⁴We were unable to purchase this tool due to its cost; however, the tool developers provided the use of this tool for this work.

⁵We were unable to purchase this tool due to its cost; the information on SHARPE in Table 1 is from the published literature and discussions with the tool developer.

technique is used to combine the models. Consequently, the model state space increases only linearly and not geometrically with the number of components modeled.

The inability of most commonly used reliability prediction tools to allow appropriate coverage values to be included in the model at any design level is also a serious problem. These prediction tools use a succinct form of system description and automatically generate the Markov model from this description. As a result, they cannot allow for the

coverages associated with each level of the error handling hierarchy to be included in those models. CARE III, CRAFTS [CRAF 88], and HARP have this problem. Actual fault-tolerant designs typically have hierarchical treatment of errors, and not all errors are necessarily treated the same way at each level. Recovery is also typically hierarchical. If the fault tolerance mechanisms associated with the individual levels in the error handling hierarchy, together with their respective coverages are not accurately included in the model, then the model will not provide a good approximation of the system, and will result in inaccurate reliability predictions. As will be shown, tools which model coverage only at the lowest level of the design may provide grossly inaccurate predictions of reliability.

SNARC (Semi-Markov Numerically Approximated Reliability Calculations) is a preliminary version of a new reliability prediction tool that overcomes these serious problems. SNARC was developed at The Aerospace Corporation to provide satellite system developers with inexpensive solutions to these problems. Large satellite systems can be modeled using this tool. Appropriate and different coverage values can be included in the model at any design level. An advantage of the numerical technique used in SNARC is that it allows large models to be solved on a PC. The other tools listed in Table I are unable to do this.

SHARPE [Sahn 86] is an existing reliability prediction tool that can also deal effectively with large state spaces and can include appropriate and different coverage values in the model at any design level. However, SHARPE has not been used in the development of any satellite systems that we are aware of, probably due to the tool's relatively high cost.

The report is organized into four sections. First, the numerical technique that forms the basis of SNARC is described. Second, the importance of accounting for coverages at all levels in the design

hierarchy is shown. Third, SNARC is used to predict the reliability of a hypothetical multicomputer system to illustrate the tool's usefulness in predicting the reliability of the complex systems currently under design. Finally, conclusions and directions for further research are described.

II. A NUMERICAL TECHNIQUE FOR SOLVING A HIERARCHY OF SEMI-MARKOV MODELS

A. DESCRIPTION OF THE TECHNIQUE

To overcome the two problems of large state spaces and allowing appropriate coverage values to be included in the model at any design level, the use of a hierarchy of semi-Markov models is employed. One can proceed downward from the highest level to the lowest level of the design, converting each level into a set of semi-Markov models, one for each separate subsystem in a design level. Each semi-Markov model contains only the information important to the subsystem (i.e., the cumulative failure rates of components and coverages) at that level in the design, suppressing unnecessary information about lower levels. Modeling the design as a hierarchy of semi-Markov models avoids the state space problem, since the number of states does not grow exponentially with the number of components, but instead only linearly. The hierarchical approach allows the coverage values associated with each subsystem's error handling mechanisms at any level to be included as parameters in the appropriate subsystem model at each level. In order to solve a model at any given level, parameters from the models at the lower level are required. Thus, a technique is needed to map correctly these parameters from the lower level models to the models at the next higher level.

The proposed technique to combine the models at different levels is to first take the same time slice from time t to $t + \Delta t$ through all levels of the hierarchy. Then, the reliability for all of the models at the lowest level is calculated. Next, the approximate increase in the cumulative failure rate (CFR) for that time slice of these same models is calculated and "inserted" as parameters in the next higher level models. This iterative procedure is followed until the reliability of the highest level model is calculated. The entire process is repeated for n time slices over the time interval that the reliability of the system is to be determined.

The reliability of the highest level model at the last time slice will be the reliability of the entire system for the given total time interval.

The approximate increase in the cumulative failure rate for each model in the interval from time t to $t + \Delta t$ is determined by calculating $[R_i(t + \Delta t) - R_i(t)]/R_i(t)$ (see Appendix). The amount of error in the approximation is a function of the size of time step Δt used in solving the model due to the integration technique used.

Figure 1 illustrates the use of this modeling approach for a system that contains three design levels. Figure 2 shows the fault tree representation of this design. Level zero (Figure 1a) represents the top of the design hierarchy, and level two (Figure 1c) represents the bottom of the design hierarchy. Each design level L_i is composed of three components, of which two must work. The model at any lower level represents a more detailed model of a component at the next higher level.

The models are combined from the lowest design level up to the highest design level for each time slice, starting from time zero and ending with the total mission time. First, the reliability of design level two (Figure 1c) at time slice (t to $t + \Delta t$) is determined by solving its semi-Markov model. Then, the increase in the level one CFR in the time interval from t to $t + \Delta t$ is calculated by


$$[R_2(t + \Delta t) - R_2(t)]/R_2(t)$$

where $R_2(t)$ is the reliability of design level two at time t . Next, the increase in the level one CFR is used to determine the reliability of design level one (Figure 1b) at time $t + \Delta t$ by solving its semi-Markov model. Then the increase in the level zero CFR is calculated for the same time interval by


$$[R_1(t + \Delta t) - R_1(t)]/R_1(t).$$

Design Level

0



1



2




Figure 1. Modeling Approach from Three Design Levels

The increase in the level zero CFR is used in the semi-Markov model for design level zero (Figure 1a) to calculate the reliability of design level zero, $R_0(t + \Delta t)$, and thus the entire system at time $t + \Delta t$.

The accuracy of this technique's final result is a function of the number of the time slices used to solve the models. The more time slices used to solve the models, the more accurate the approximation. As an initial guideline to finding the minimum number of time slices that results in an accurate solution, the tool user should first solve the models using seventy time slices per year. Then, for each repetition, the user should increase the number of time slices to solve the models by 30 time slices per year until the model solution is stable in the decimal place of desired accuracy (e.g. no change takes place in the fifth decimal place after three or four successive model solutions).

The need to perform repeated solutions to find the minimum number of time slices to ensure an accurate solution is not a permanent limitation. It is only a consequence of the tool's current method of implementation. The inclusion of other solution techniques would remove this limitation.

B. CURRENT LIMITATIONS OF THE TECHNIQUE

The limitations discussed in this section are common to reliability prediction tools in general and specifically to all of the tools listed in Table I. It is proposed to solve these problems for SNARC, thereby advancing the state-of-the-art in reliability prediction tools even further than the advances already made by SNARC.

First, when using time-varying failure rates, all the tools in Table I, except CRAFTS (which is unable to model time-varying failure rates altogether), must model all standby spares as "hot." Fortunately, this limitation results in conservative reliability estimates. However, the magnitude of the inaccuracy cannot be predicted in general because the

reliability estimates are a function of the various failure rates involved and the number of spares present at each level of the design hierarchy. We believe that this problem can be solved in SNARC with a modified but straightforward solution technique.

Second, none of the tools in Table I are able to deal accurately with transients that propagate up through the design hierarchy. In CRAFTS, CARE III, and HARP, if the fault and error handling mechanisms, located at the lowest level in the hierarchy, cannot detect and recover from a fault's effects, then the system fails, even though in the actual design, higher level mechanisms may be included to recover adequately from the fault's effects. SNARC and SHARPE do better and are able to deal with permanent faults that propagate up the levels of the hierarchy. Unfortunately, the current versions of both SNARC and SHARPE are only able to deal with transient fault effects that propagate up the hierarchy by treating them as permanent. We propose to solve this problem in SNARC, but it will require additional funding and time.

III. IMPORTANCE OF COVERAGES IN THE DESIGN HIERARCHY

In this section, we use SNARC to show that extremely inaccurate predictions of design reliability are produced by tools that cannot model coverages throughout the entire design hierarchy. To show this, we take an example system and model it with analytical calculations, SNARC, and CRAFTS using different coverage values throughout the design hierarchy. CRAFTS is a reliability prediction tool based on ARIES [Ng 80, Maka 82]. We used the analytical calculations as the reference in comparing results. CRAFTS was used as a representative of the tools which are unable to model the effects of coverages throughout the design hierarchy. The results of CRAFTS were compared to the analytical and the SNARC results to determine how far the results of this class of tools are removed from the correct values (i.e., the results of the analytical calculations and SNARC).

The system in Figure 1 was used as the example. Initially, the reliability of the system was calculated using perfect coverage to provide a comparative baseline for the three techniques used. The failure rate of the components at level 2 was defined as 1×10^{-6} failures/hour. Table 2 presents the results of (1) analytical calculations, (2) SNARC, and (3) CRAFTS.

Tables 3 through 5 present the reliability results for the system in Figure 1, using coverage values of 0.99, 0.95, and 0.90 respectively at each design level. Each table presents the results from analytical calculations, SNARC, and CRAFTS. The slight difference between some of the analytical calculations and the SNARC results are due to errors introduced by the simple integration technique used in SNARC. The CRAFTS models used for design levels zero and one were reliability block diagrams, since constructing a Markov model of the system for use in CRAFTS is precluded by the large state space of the model, which is greater than 64,000 states. The results from CRAFTS show the impact of

Table 2. Triad System Reliability
Coverage = 1

Year	Analytical Calculations	SNARC	CRAFTS
1	1.000000	1.000000	1.000000
2	1.000000	1.000000	1.000000
3	1.000000	1.000000	1.000000
4	1.000000	1.000000	1.000000
5	1.000000	1.000000	1.000000
6	1.000000	1.000000	1.000000
7	1.000000	1.000000	1.000000
8	0.999999	0.999999	0.999999
9	0.999998	0.999998	0.999998
10	0.999996	0.999996	0.999996

Table 3. Triad System Reliability
Coverage = 0.99

Year	Analytical Calculations	SNARC	CRAFTS
1	0.999999	1.000000	1.000000
2	0.999999	0.999999	1.000000
3	0.999997	0.999997	1.000000
4	0.999994	0.999994	1.000000
5	0.999990	0.999990	1.000000
6	0.999984	0.999984	1.000000
7	0.999976	0.999976	0.999999
8	0.999964	0.999964	0.999999
9	0.999947	0.999948	0.999997
10	0.999925	0.999925	0.999994

Table 4. Triad System Reliability
Coverage 0.95

Year	Analytical Calculations	SNARC	CRAFTS
1	0.999965	0.999965	1.000000
2	0.999918	0.999918	1.000000
3	0.999857	0.999857	1.000000
4	0.999780	0.999780	1.000000
5	0.999685	0.999685	1.000000
6	0.999568	0.999569	0.999999
7	0.999427	0.999428	0.999997
8	0.999257	0.999258	0.999994
9	0.999053	0.999055	0.999988
10	0.998811	0.998813	0.999978

Table 5. Triad System Reliability
Coverage = 0.90

Year	Analytical Calculations	SNARC	CRAFTS
1	0.999741	0.999741	1.000000
2	0.999433	0.999433	1.000000
3	0.999072	0.999073	1.000000
4	0.998652	0.998653	0.999999
5	0.998166	0.998186	0.999998
6	0.997609	0.997612	0.999994
7	0.996972	0.996975	0.999988
8	0.996247	0.996251	0.999977
9	0.995425	0.995430	0.999958
10	0.994496	0.994502	0.999928

including coverage at design level two only. No coverage values could be associated with the CRAFTS models at design levels zero and one, since the coverages associated with those levels could not be included in the reliability block diagram models. (Neither can the coverages at the higher levels be represented by use of the fault tree notation implemented by other reliability prediction tools.) Of course, it could be claimed that an estimated lower coverage value should have been used at design level two for CRAFTS to approximate the coverage at the higher levels. However, this is not an adequate way to represent the coverages associated with these mechanisms. Any coverage value used in this manner for modeling would have been sheer conjecture, since this parameter would represent an abstraction of the coverages associated with several levels of state-dependent error handling mechanisms into a single value. (The mechanisms that implement these coverages are not enabled by the failure of every lower level component.)

The ability to include in the models the coverages present at the higher design levels significantly affects the reliability prediction results. Tables 2 through 5 show that the results from SNARC agree very closely with the reference results, but that the results from CRAFTS differed by a significant amount. Table 6 shows the unreliability of each of the SNARC models with imperfect coverage divided by the unreliability of the equivalent CRAFTS models. When coverage is 0.9, the SNARC results show an unreliability 76 times greater than the CRAFTS models. Thus, the coverages associated with the higher design levels must be included in the reliability prediction models to allow designers to assess more accurately the impact of these coverages on design reliability.

Table 6. Ratio of Model Unreliability for Triad

Coverage	Ratio
0.99	12.50
0.95	53.95
0.90	76.36

In reviewing the modeling performed using SNARC on the system shown in Figure 1, we notice that several benefits have accrued. A significant reduction in the state space has been achieved in the model of this system. The "flat" model (i.e., a single Markov model of the type created by HARP) of the system would have more than 64,000 states. The equivalent model using the hierarchy of semi-Markov models requires only nine states. Also, the impact of the coverages associated with each design level has been included in the SNARC models.

IV. PREDICTING THE RELIABILITY OF A SYSTEM DESIGN

In this section, SNARC is used to model a hypothetical but representative multicomputer to illustrate the usefulness of the numerical technique in predicting the reliability of the complex systems under design today.

The system consists of ten computers, eight of which must work for the system to function. Each computer is composed of a memory subsystem, a processor subsystem, and an I/O subsystem. The network between computers was not included merely for convenience. Figures 3 and 4 illustrate the block diagram and the fault tree for one of the ten computers.

Each computer's memory subsystem consists of three memory modules, of which two must work. A memory module consists of a redundant 41-chip memory array and a nonredundant 2-chip interface. Thirty-nine of the memory array chips must work and two are spares. All chips have a failure rate of 1×10^{-7} failures/hr. The coverage of memory chip failures is 0.998, and the interface has a coverage of 0.

Each computer's processor subsystem consists of three pairs of self-checking CPUs and a nonredundant 2-chip recovery processor. At least two of the three CPUs must work. Each pair of processors consists of six chips and is replaced as a unit. The failure rate of each chip is 1×10^{-7} failures/hr. The coverage for processor recovery is 0.995.

A computer's I/O subsystem consists of standby redundant ports. Each port has a failure rate of 6×10^{-7} failures/hr. The coverage for the I/O subsystem is 0.99.

Table 7 presents the reliability of this system after 10 yr, and shows the results of CRAFTS and SNARC. SNARC was able to solve this model on a PC in 30 sec. Additionally, the SNARC models were evaluated

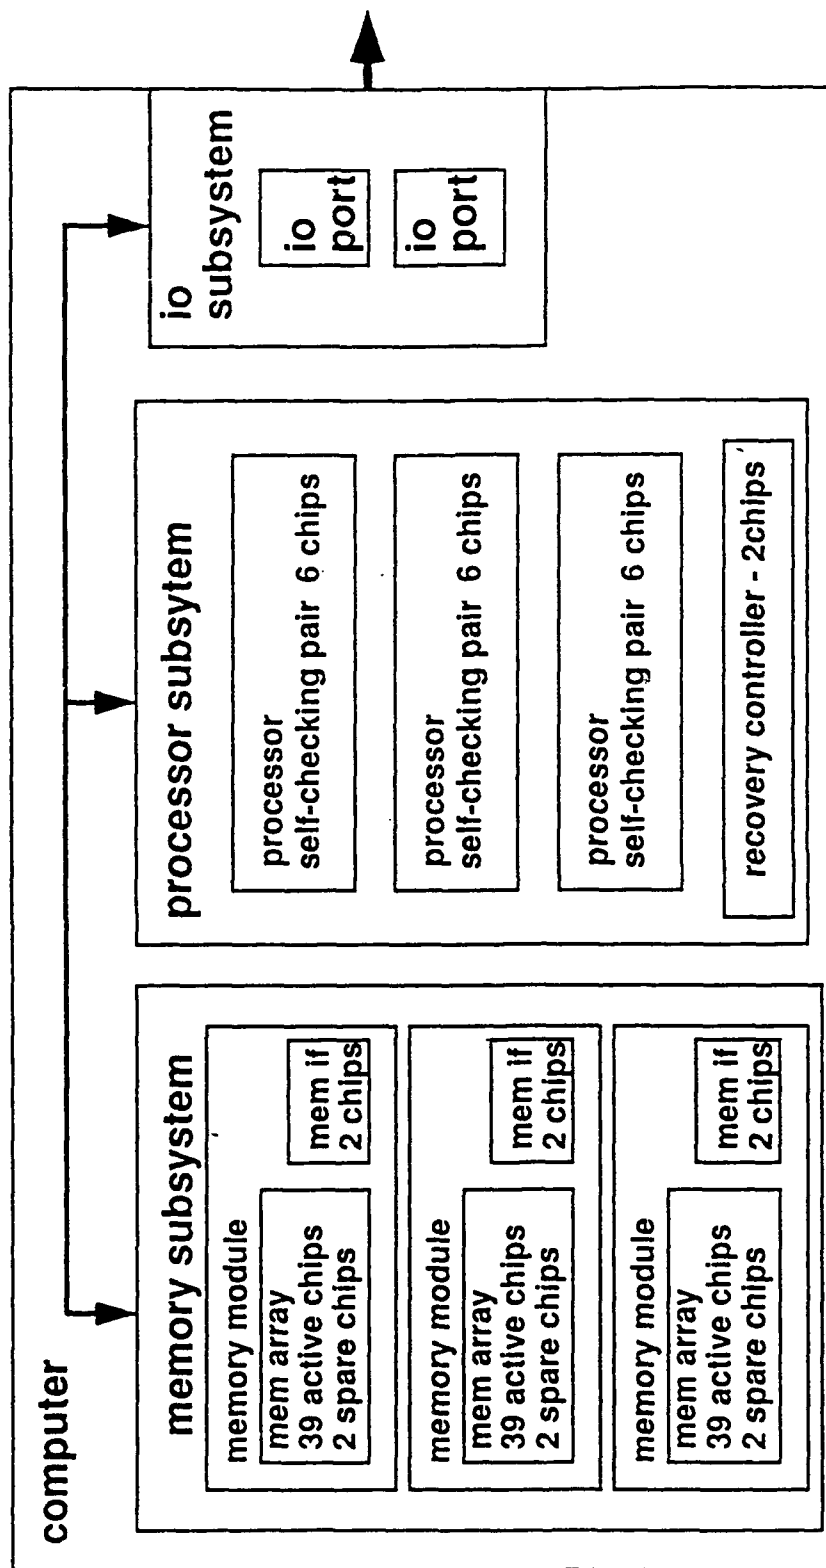


Figure 3. Organization of One Computer in SNARC Modeling

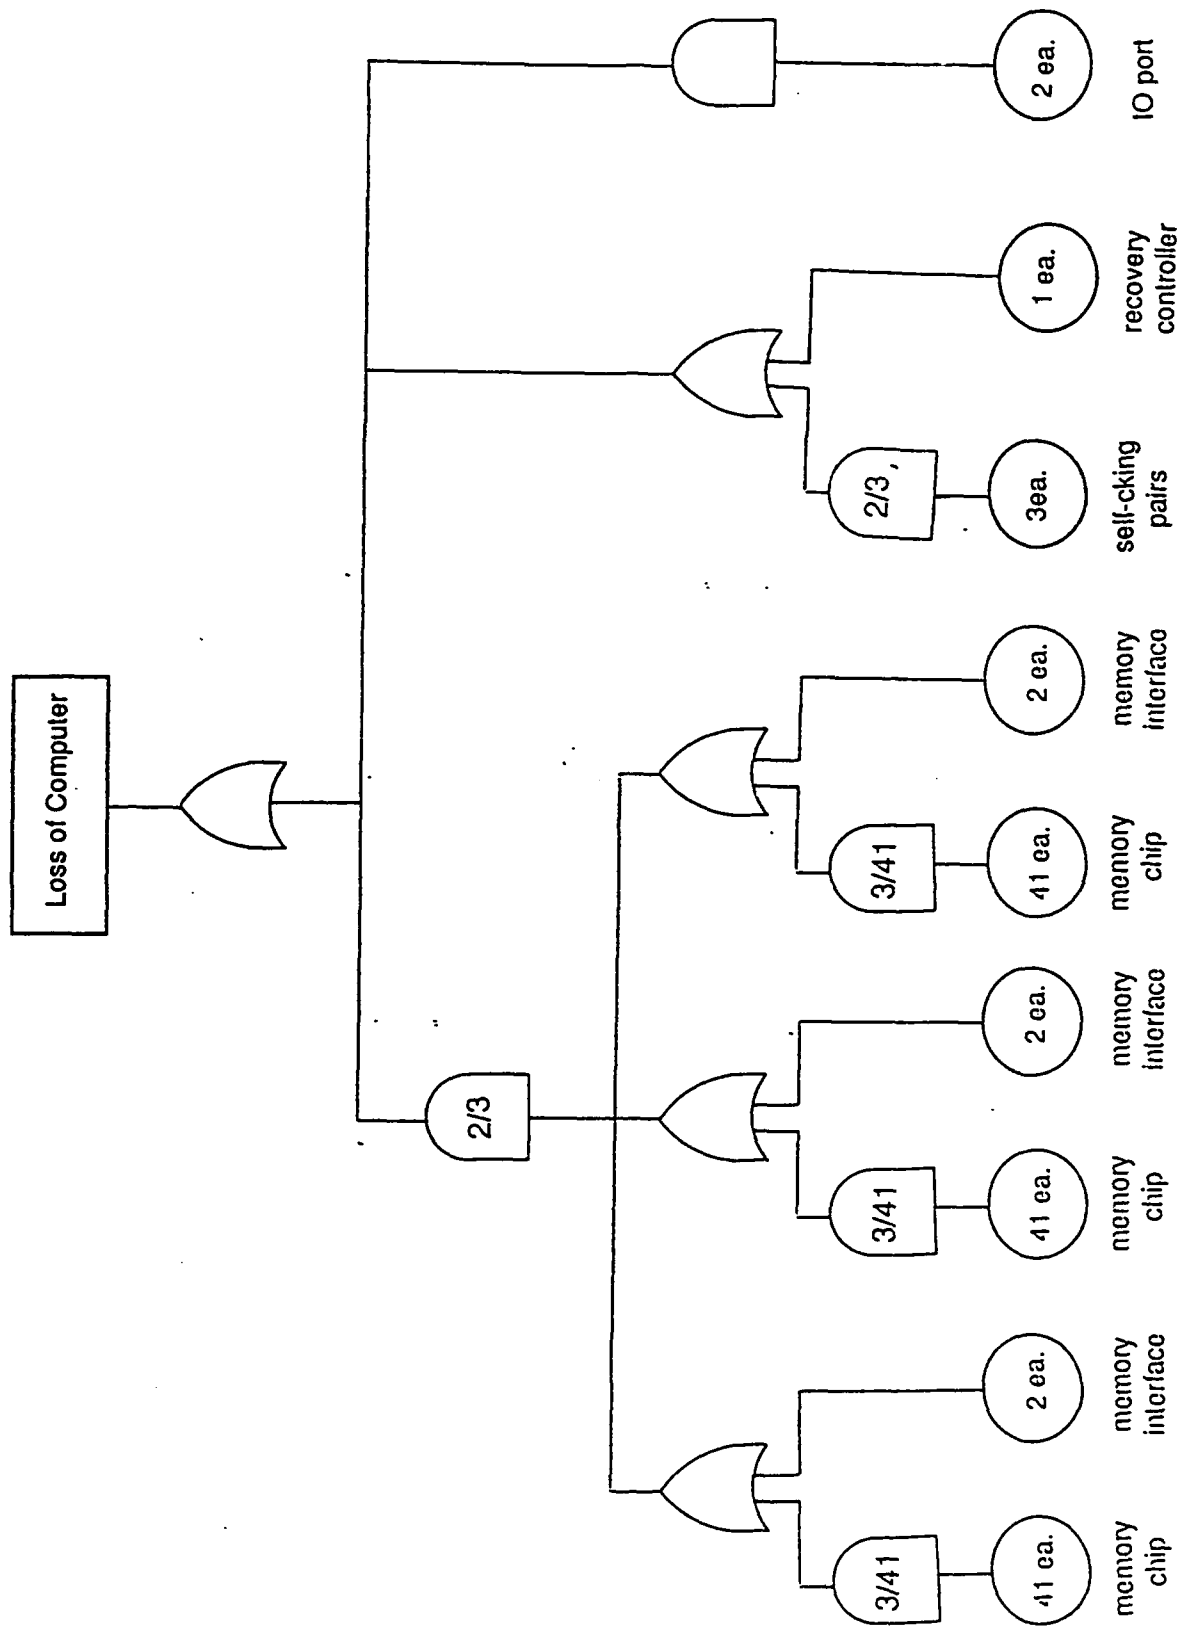


Figure 4. Fault Tree for Computer in SNARC Modeling

Table 7. Multicomputer Reliability

Year	CRAFT Results	SNARC Results			
	Coverage	Coverage			
	1.00	1.00	0.99	0.95	0.90
1	0.999999	0.999999	0.999788	0.998837	0.997403
2	0.999990	0.999991	0.999546	0.997550	0.994570
3	0.999963	0.999963	0.999261	0.996128	0.991482
4	0.999900	0.999900	0.998918	0.994550	0.988115
5	0.999780	0.999780	0.998493	0.992791	0.984437
6	0.999574	0.999575	0.997958	0.990819	0.980413
7	0.999247	0.999248	0.997276	0.988593	0.975997
8	0.998756	0.998758	0.996403	0.986064	0.971135
9	0.998051	0.998054	0.995287	0.983176	0.967680
10	0.997071	0.997074	0.993865	0.979860	0.959823

where the coverages, at the computer memory subsystem and multicomputer levels, were 0.99, 0.95, and 0.9. However, CRAFTS was unable to reflect the coverage values at those design levels. Again, the inclusion of coverage values at higher design levels has significantly affected the reliability prediction of the design. Table 8 shows the unreliability of the SNARC models with imperfect coverage at the memory subsystem and the multicomputer levels, divided by the CRAFTS results with perfect coverage at those same levels. When the coverage at those two levels is 0.9, the SNARC results show an unreliability almost 14 times greater than the CRAFTS results. We were unable to use CARE III and HARP for comparison, since we were unable to get CARE III and HARP to model this system, which resulted from either exceeding the tools' capabilities or the computing resources on which the tools ran.

Monte Carlo simulations of the same multicomputer were run to validate that the numerical technique of SNARC was generating the correct results for all the models. Tables 9, 10, 11, and 12 show the results of SNARC and the Monte Carlo simulations. In terms of unreliability, the Monte Carlo and SNARC results differ by only approximately 5%. The Monte Carlo results provide support that the numerical technique, as implemented in SNARC, generates the correct solutions.

Table 8. Ratio of Model Unreliability for Multicomputer

Coverage	Ratio
0.99	2.09
0.95	6.88
0.90	13.72

Table 9. Multicomputer Reliability
Coverage = 1

Year	Monte Carlo	SNARC
1	1.000000	0.999999
2	0.999983	0.999991
3	0.999983	0.999963
4	0.999950	0.999900
5	0.999867	0.999780
6	0.99617	0.999575
7	0.999333	0.999248
8	0.998733	0.998758
9	0.997967	0.998054
10	0.997300	0.997074

Table 10. Multicomputer Reliability
Coverage = 0.99

Year	Monte Carlo	SNARC
1	0.999900	0.999788
2	0.999650	0.999546
3	0.999317	0.999261
4	0.999017	0.998918
5	0.998550	0.998493
6	0.998150	0.997958
7	0.997600	0.997276
8	0.996533	0.996403
9	0.995167	0.995287
10	0.993783	0.993865

Table 11. Multicomputer Reliability
Coverage = 0.95

Year	Monte Carlo	SNARC
1	0.998717	0.998837
2	0.997517	0.997550
3	0.996150	0.996128
4	0.994733	0.994550
5	0.993017	0.992791
6	0.990867	0.990819
7	0.986667	0.988593
8	0.986700	0.986064
9	0.983867	0.983176
10	0.980717	0.979860

Table 12. Multicomputer Reliability
Coverage = 0.90

Year	Monte Carlo	SNARC
1	0.997383	0.997403
2	0.994733	0.994570
3	0.991400	0.991482
4	0.987867	0.988115
5	0.983800	0.984437
6	0.979917	0.980413
7	0.975767	0.975997
8	0.971333	0.971135
9	0.965533	0.967680
10	0.960300	0.959823

V. CONCLUSIONS AND FUTURE WORK

We have presented a numerical technique for solving a hierarchy of semi-Markov models to analyze the reliability of large, closed, fault-tolerant, on-board computer systems. A preliminary tool called SNARC was developed using this approach, with the intent of providing all system developers with an inexpensive way to treat the coverages associated with each level of the design hierarchy in a more accurate manner. In addition, the approach used by the tool avoids the large state space problem. We used SNARC to solve the semi-Markov models of several example systems (that have large corresponding Markov models) on a PC, in a short period of time (e.g., 30 sec or less). We were able to show the significant sensitivity of design reliability to the effects of imperfect coverage at the various levels of the design hierarchy, since we were able to include the effects of these coverages in the SNARC models. The tools CRAFTS, HARP, and CARE III could not do this.

Future work is proposed in three areas. First, three modifications are proposed to make the SNARC program easier for designers to use. These modifications are listed in the order of their importance. In order for designers to use the tool, documentation describing the tool and how to use it needs to be written. A numerical technique should be identified and included in SNARC that removes the need for the designer to find, in an interactive manner, the minimum number of time steps that results in an accurate solution. Also, a graphical interface would allow users to input models easily.

Second, four modifications to expand SNARC's capabilities are proposed, in order of their importance. Incorporate techniques to correctly model the use of "warm" spares throughout the design hierarchy. Develop techniques that allow the hierarchical modeling of transients. Add the capability to model near-coincident faults throughout the

hierarchy. Include different Fault/Error Handling Models (FEHMs) in SNARC to enable users to select the FEHM that approximates best the fault and error handling mechanisms at each level of the design hierarchy.

Finally, continued testing of SNARC needs to be done to ensure that the tool is "bug" free.

BIBLIOGRAPHY

1. [Bavu 84a] S. Bavuso, J. Brunelle, and P. Peterson,
"CARE III Hands-on Demonstration and Tutorial,
"NASA Technical Memorandum 85811, May 1984.
2. [Bavu 84b] S. Bavuso, P. Peterson, and D. Rose,
"CARE III Model Overview and User's Guide,
"NASA Technical Memorandum 85810, June 1984.
3. [CRAF 88] CRAFTS User Reference and Manual, Jan. 1988.
4. [Duga 86] J. Dugan, K. Trivedi, M. Smotherman, and R. Geist,
"The Hybrid Automated Reliability Predictor,"
AIAA Journal of Guidance, Control, and Dynamics,
May-June 1986, pp. 319- 331.
5. [Duga 89] J. Dugan, K. Trivedi,
"Coverage Modeling for Dependability Analysis of Fault-
Tolerant Systems," *IEEE Transactions on Computers*,
Vol. 38, No. 6, June 1989, pp. 775-787.
6. [HARP 86] "HARP: The Hybrid Automated Reliability Predictor
Introduction and User's Guide," NASA Langley Research
Center, Sept., 1986.
7. [Maka 82] S. Makam and A. Avizienis,
"ARIES 81: A Reliability and Life-Cycle Evaluation Tool
for Fault-Tolerant Systems," *Proceedings IEEE 12th
Fault-Tolerant Computing Symposium*, June 1982, pp.
267-274.

8. [Ng 80] Y. Ng and A. Avizienis,
"A Unified Reliability Model for Fault-Tolerant
Computers," *IEEE Transactions on Computers*, Vol. C-29,
No. 11, Nov. 1980, pp. 1002-1011.
9. [Sahn 86] R. Sahner and K. Trivedi,
"A Hierarchical Combinatorial-Markov Method for Solving
Complex Reliability Models," *Proceedings ACM/IEEE
Fall Joint Computing Conference*, Nov. 1986.
10. [Triv 81] K. Trivedi and R. Geist,
"A Tutorial on the CARE III Approach to Reliability
Modeling," NASA Contractor Report 3488, Dec. 1981.
11. [Triv 82] K. Trivedi,
Probability and Statistics with Reliability, Queueing
and Computer Science Applications, Prentice-Hall,
Englewood Cliffs, NJ, 1982.
12. [Triv 85] K. Trivedi, R. Geist, M. Smotherman, and J. Dugan,
"Hybrid Modeling of Fault-Tolerant Systems,"
*Computers and Electrical Engineering, An International
Journal*, vol. 11, no. 2 & 3, 1985, pp. 87- 108.

APPENDIX. APPROXIMATING THE INCREASE IN $H_i(t)$ FROM TIME t TO $t + \Delta t$

We know that the cumulative failure rate $H_i(t) =$

$$\int_0^t h_i(\tau) d\tau = \int_0^t -R_i'(\tau)/R_i(\tau) d\tau \text{ [Triv 82]}$$

$$\begin{aligned} \text{The integral } \int_0^t -R_i'(\tau)/R_i(\tau) d\tau &= \int_0^{\Delta t} -R_i'(\tau)/R_i(\tau) d\tau \\ &+ \int_{\Delta t}^{2\Delta t} -R_i'(\tau)/R_i(\tau) d\tau + \int_{2\Delta t}^{3\Delta t} -R_i'(\tau)/R_i(\tau) d\tau + \dots + \int_{(N-1)\Delta t}^{N\Delta t} -R_i'(\tau)/R_i(\tau) d\tau, \end{aligned}$$

where $N\Delta t = t$.

We can approximate $\int_t^{t+\Delta t} [R_i(\tau + \Delta\tau) - R_i(\tau)]/[\Delta\tau R_i(\tau)] d\tau$ by $[R_i(t + \Delta t) - R_i(t)]/R_i(t)$, when

Δt is very close to zero. Substituting into the above equations, we obtain

$$\begin{aligned} H_i(t) &= \int_0^t h_i(\tau) d\tau = [R_i(0 + \Delta t) - R_i(\Delta t)]/R_i(0) + [R_i(0 + 2\Delta t) - R_i(0 + \Delta t)]/R_i(0 + \Delta t) + \dots \\ &+ [R_i(0 + (N\Delta t)) - R_i(0 + (N-1)\Delta t)]/R_i(0 + (N-1)\Delta t) \end{aligned}$$

Therefore, we can approximate the increase in the cumulative failure rate $H_i(t)$, from time t to $t + \Delta t$ by determining $[R_i(t + \Delta t) - R_i(t)]/R_i(t)$, and we can approximate $H_i(t)$ by summing all incremental increases in $H_i(t)$ from time 0 to the end of mission life time.